

BINARY DECISION TREE-BASED ARBITRATOR FOR PACKETIZED
COMMUNICATIONS

FIELD OF THE INVENTION

5 The present invention relates in general to communication networks and systems employed for the transport of digital telecommunication signals, and is particularly directed to a new and improved binary decision tree-based arbitration scheme that is executable
10 by a supervisory control processor of a time division multiplex (TDM)-based communication system, and is operative to select the next packet to be transmitted from a plurality of virtual circuits, any number of which may have one or more packets awaiting transmission over
15 a serialized digital communication link.

BACKGROUND OF THE INVENTION

Digital communication systems enable telecommunication service providers (for example, a competitive local exchange carrier (CLEC), such as an
20 internet service provider (ISP)), to provide various types of high speed digital service over network circuits of an incumbent local exchange carrier (ILEC), such as a Bell operating company (RBOC), serving a number of customer premises equipments (CPEs) having a wide range
25 of operational bandwidths and digital subscriber line termination capabilities. A reduced complexity example of such a digital communication network architecture is

diagrammatically illustrated in Figure 1 as comprising a PCM communication link (such as an optical fiber) 10, through which a network (cloud) 20 at a 'west' end of the link 10 may transmit and receive digital telecommunication signals (e.g., packetized T3 traffic) with respect to customer premises equipments (CPEs) served by a remote termination site (RTS) 30 at an 'east' end of the PCM link 10.

As shown in Figure 2, in order to handle transmission requests from the virtual circuits (VCs) that are associated with various customer equipments (CPEs), the equipment shelf of a remote terminal 30 terminating the 'east' end of the PCM link 10 typically contains a supervisory communication control processor 40, one of the functions of which is to arbitrate/grant the interfacing of packets awaiting transmission from one or more virtual circuit ports $P_{VC1} - P_{VCN}$ to the PCM serial link 10. Now although a maximal speed/memory processor architecture could be used to execute a sequential polling scheme to accomplish this task, doing so would not only be cost-prohibitive from a commercial application standpoint, but could be expected to be a less than efficient use of available bandwidth (e.g., checking each and every virtual circuit port including those having no data to transmit).

SUMMARY OF THE INVENTION

In accordance with the present invention, this virtual circuit servicing problem is successfully remedied by employing a binary decision tree-based transmission assignment mechanism, that is operative to point directly to the virtual circuit currently having the highest transmission priority as a result of the most immediately precedent transmission, and to update the transmission priority of the nodes of the decision tree each time a packet awaiting transmission has been serviced.

The transmission priority scheme of the invention is comprised of $N+1$ sets of nodes containing $2^{N+1}-1$ nodes that define a binary decision tree-structure. A respective i th set of nodes comprises 2^{i-1} nodes, wherein i is greater than or equal to 1, and less than or equal to $N+1$. The nodes of a given set are connected to those of an adjacent set by binary-split branches. Thus, a first node set contains a single (or root) termination node, associated with a port associated with a serialized digital communication link. The second node set contains two nodes that are coupled by a pair of respective branches to the root node, and so on, to the $N+1$ th set of nodes at the end of the decision tree opposite to the root node, wherein an $(N+1)$ th node set contains $2^{(N)}$ leaf nodes associated with respective communication ports of 2^N virtual circuits.

For each of the 2^N leaf nodes of the decision tree, information is stored representative of the transmission priority of a packet that may be awaiting transmission from its associated communication port. In addition, 5 associated with each of those nodes that branch to respective pairs of downstream nodes (toward the virtual circuit ports) is a respective stored code value, or 'pointer' code that, points to whichever one of its two branched nodes is associated with a higher packet 10 transmission priority.

As a result, as one traverses the decision tree along sequentially split branch paths from the highest priority leaf node toward the root node, the pointer code of the next immediately adjacent upstream node will 15 always point to the node-branch path leading to the highest priority leaf node. This means that the root node will always point to the highest priority leaf node. Namely, the highest priority transmission priority code will 'ripple' through the decision tree to the pointer 20 data stored at the root node. This means that in order to determine from which virtual circuit the next packet is to be obtained and coupled to the serialized data link, the communication controller need look no further than the root node.

25 Once the current packet request has been honored (a virtual circuit delivers a packet to the transmitter (the root of the decision tree)) it recalculates its priority, in preparation for the next decision tree operation.

Recalculating the priority on every packet transmission avoids imparting additional delay to high-priority, time-sensitive, traffic flowing in a congested network. This priority recalculation may include one or more of a 5 number of factors, such as, but not limited to arrival time of the next packet in the queue, the current bandwidth utilization of the virtual circuit, the type of the next packet in the queue, and user-assigned weight or quality of service.

10 With priority recalculated, the pointer codes of the decision tree are updated to reflect the leaf node which now points to the highest priority port (which could be the same as or different from the previous port). The communications controller need only look to the root node 15 for the pointer to the leaf node and thereby virtual circuit port from which the next packet is to be selected for transmission.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a reduced complexity diagram of a 20 digital communication network architecture;

Figure 2 shows a supervisory control processor employed by a remote unit terminating a PCM link to arbitrate/assign the interfacing of packets awaiting transmission from one or more virtual circuits served by 25 the remote unit; and

Figure 3 diagrammatically illustrates the binary decision tree-based transmission assignment mechanism of the present invention.

DETAILED DESCRIPTION

5 Before describing in detail the new and improved binary decision tree-based transmission arbitration mechanism in accordance with the present invention, it should be observed that the invention resides primarily in what is effectively a prescribed communication control
10 mechanism that is executable by the hardware and software of supervisory communications control components of conventional digital communication circuitry, including digital signal processing components and attendant supervisory control circuitry therefor, that controls the
15 operations of such circuits and components.

As a consequence, the configuration of such circuits and components and the manner in which they are interfaced with other communication system equipment have, for the most part, been illustrated in the drawings by readily understandable block diagrams, which show only those specific details that are pertinent to the present invention, so as not to obscure the present disclosure with details which will be readily apparent to those skilled in the art having the benefit of the description
20 herein. Thus, the diagrammatic illustrations are primarily intended to show the major components and functional operations of the invention in the context of
25

a present day digital communication network in a convenient functional grouping, whereby the present invention may be more readily understood.

Attention is now directed to Figure 3, which 5 diagrammatically illustrates a reduced complexity example of the binary decision tree-based transmission assignment mechanism of the present invention. As pointed out above, this transmission assignment mechanism is executable by a supervisory communications control processor of remote 10 communications unit terminating a serialized digital communication link, such as that shown in Figures 1 and 2, described above, to expeditiously arbitrate/assign the interfacing of packets awaiting transmission from one or more of a plurality of virtual circuits served by the 15 remote unit.

The reduced complexity example of Figure 3 illustrates the case of arbitrating packet transmission requests for up to $2_{(N=3)}$ or eight virtual circuits VC1 - VC8. It is to be understood, however, that the invention 20 is not limited to this or any particular number. The choice of eight virtual circuits in the example of Figure 3 serves to reduce the complexity of the drawings, while still providing a useful illustration of the principles of the invention.

25 Irrespective of the number of virtual circuits, the transmission assignment network is comprised of $N+1$ sets of nodes containing $2^{N+1}-1$ nodes that define a binary decision tree-structure. In the present example of

servicing eight virtual circuits, the decision tree-structure of Figure 3 is comprised of $(N=3)+1$ or four sets of nodes (comprised of node sets S1, S2, S3 and S4), that contain a total of fifteen nodes N1 - N15, that are 5 interconnected by branches B of the decision tree. A respective ith set of nodes comprises 2^{i-1} nodes, wherein i is greater than or equal to 1, and less than or equal to N+1.

Thus, a first ($i=1$) node set S1 contains a single 10 (or root) termination node N1, that is associated with a port P1 coupled to the serialized digital communication link 10. As 2^N is equal to or greater than the number of virtual circuit communication ports, at the other end of the decision tree, a fourth ($i=4$) node set S4 of nodes 15 contains $2^{(N=3)}=8$ leaf nodes N8-N15 that are associated with respective communication ports P2-P9 for up to eight virtual circuits VC1-VC8. The ($i=2$) second node set S2 contains two nodes N2 and N3, coupled by respective branches B12, B13 to the root node N1. The third ($i=3$) 20 node set S3 contains four nodes N4 - N7, that are coupled by respective branches B24, B25; B36, B37 to nodes N2 and N3 of the second set S2. The eight leaf nodes N8 - N15 of the fourth node set S4 are coupled by respective branches B48, B49; B510, B511; B612, B613; and B714, B715 to nodes 25 N4-N7 of the third set S3.

For each of the 2^N leaf nodes of the decision tree, there is first stored information (a transmission priority code shown by a transmission priority code box

TPC) representative of the priority of transmission of a packet that may be awaiting transmission from its associated communication port. In addition, associated with each of those nodes that branch to respective pairs 5 of downstream nodes (toward the virtual circuit ports), namely with nodes N1-N7 (as none of the leaf nodes N8-N15 branches to a further pair of nodes), is a respective stored code value or 'pointer' code (denoted by box PNTR), that points to whichever one of its two branched 10 nodes is associated with a higher packet transmission priority.

For purposes of the present example, let it be assumed that of the eight ports P2-P9, virtual circuit VC5 (associated with port P6) has the highest 15 transmission priority. Thus, the pointer code associated with immediately upstream branched node N6 of the third set S3 of nodes (shown by arrow 6-12) will point to leaf node N12, corresponding to virtual circuit port P6, which has a higher priority than the other lead node N13 20 branched from node N6. Since the pointer associated with branched node N6 of the third set S3 of nodes points to the leaf node (N12) having the highest transmission priority, then the pointer (shown by arrow 3-6) associated with its immediately upstream branched node N3 25 of the second node set S2 will point to node N6 and thereby to leaf node N12, whose transmission priority is higher than either of the leaf nodes (N14, N15) branched

from the other node N7 which, in turn, is branched from node N3.

In a similar manner, as the pointer (3-6) associated with branched node N3 of the second set S2 of nodes 5 points to the leaf node (N12) having the highest transmission priority, then the pointer (shown by arrow 1-3) associated with its immediately upstream branched node (here root node N1) of the first node set S1 will point to node N3 and thereby to leaf node N12. Since node 10 N3 points to a (highest priority) leaf node (N12) whose transmission priority is higher than any of the leaf nodes pointed to by the other node (N2) branched from root node N1, the root node N1 necessarily points the highest priority leaf node N12.

15 Namely, the highest priority transmission priority code (in the present example, associated with virtual circuit VC5) will 'ripple' through the decision tree to the pointer data stored at the root node N1. This means that in order to determine from which virtual circuit the 20 next packet is to be obtained and coupled to the serialized data link, the communication controller need look no further than the root node N1, which always points to the highest priority leaf node (leaf node N12 in the present example).

25 Once the current packet request has been honored (a virtual circuit delivers a packet to the transmitter (the root of the decision tree)) it recalculates its priority, in preparation for the next decision tree operation. As

pointed out above, recalculating the priority on every packet transmission avoids imparting additional delay to high-priority, time-sensitive, traffic flowing in a congested network. This priority recalculation may 5 include one or more of a number of factors, such as, but not limited to arrival time of the next packet in the queue, the current bandwidth utilization of the virtual circuit, the type of the next packet in the queue, and user-assigned weight or quality of service. Once priority 10 has been recalculated, the pointer codes of the decision tree are updated to reflect whichever leaf node now points to the highest priority port (which could be the same as or different from the previous port). Again, the communications controller need look to only the root node 15 for the pointer to the leaf node and thereby virtual circuit port from which the next packet is to be selected for transmission.

As will be appreciated from the foregoing description, the binary decision tree-based transmission 20 assignment mechanism of the present invention provides a reduced complexity and cost scheme for servicing a large number of virtual circuits, and thereby obviates the need for a maximal speed/memory processor architecture to sequentially poll the circuits. As noted previously, not only 25 is polling a large number of circuits cost-prohibitive, but is typically a less than efficient use of available bandwidth, as it involves checking each virtual circuit port including those having no data to transmit.

While we have shown and described an embodiment in accordance with the present invention, it is to be understood that the same is not limited thereto but is susceptible to numerous changes and modifications as known to a person skilled in the art, and we therefore do not wish to be limited to the details shown and described herein, but intend to cover all such changes and modifications as are obvious to one of ordinary skill in the art.